

The fractal structure of cellular automata on Abelian groups

JOHANNES GÜTSCHOW¹, VINCENT NESME^{2*},
REINHARD F. WERNER¹

¹ *Institut für Theoretische Physik, Universität Hannover*

² *Quantum information theory, Universität Potsdam*

It is well-known that the spacetime diagrams of some cellular automata have a fractal structure: for instance Pascal's triangle modulo 2 generates a Sierpinski triangle. Explaining the fractal structure of the spacetime diagrams of cellular automata is a much explored topic, but virtually all of the results revolve around a special class of automata, whose typical features include irreversibility, an alphabet with a ring structure, a global evolution that is a ring homomorphism, and a property known as (weakly) p -Fermat. The class of automata that we study in this article has none of these properties. Their cell structure is weaker, as it does not come with a multiplication, and they are far from being p -Fermat, even weakly. However, they do produce fractal spacetime diagrams, and we explain why and how.

INTRODUCTION

The fractal structure of cellular automata (CAs) has been a topic of interest for several decades. In many works on linear CAs, the authors present ways to calculate the fractal dimension or to predict the state of an arbitrary cell at an arbitrary time step, with much lower complexity than by running the CA step by step; However, their notions

* email: vnesme@gmail.com

of linearity are quite different. Often only CAs that use states in \mathbb{Z}_p are studied.* Other approaches are more general, but still make certain assumptions on the time evolution or the underlying structure of the CA. In this work we try to loosen these restrictions as far as possible. We consider one-dimensional linear CAs whose alphabet is an abelian group. We show how they can be described by $n \times n$ matrices with polynomial entries and use this description to derive a recursion relation for the iterations of the CA. This recursion relation enables us to formulate the evolution of the spacetime diagram as a matrix substitution system, which in turn gives us the means to calculate the fractal dimension of the spacetime diagram.

Most of the methods we employ are commonly used in the study of CAs. To prove that the spacetime diagram converges we use elementary graph theory and study the graph associated with the matrix substitution system of the CA. The matrix substitution system is obtained by a recursion formula for the iterations of the CA and grouping of cells.

Our interest in the fractal structure of CAs on abelian groups stems from our study of Clifford quantum cellular automata (CQCA) [12]. We first noticed the self-similar structure while studying their long time behaviour [4, 3]. A CQCA maps Pauli matrices to tensor products of Pauli matrices times a phase. If we neglect the phase, we can identify the Pauli matrices X, Y, Z with the elements of \mathbb{Z}_2^2 via the mapping $X \mapsto \begin{pmatrix} 1 \\ 0 \end{pmatrix}, Y \mapsto \begin{pmatrix} 1 \\ 1 \end{pmatrix}, Z \mapsto \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \mathbb{1} \mapsto \begin{pmatrix} 0 \\ 0 \end{pmatrix}$. Using this mapping we can simulate CQCAs with linear CAs on the alphabet \mathbb{Z}_2^2 . The CAs corresponding to CQCAs have to fulfill rather strong conditions: they have to be reversible and preserve a symplectic form which encodes the commutation relations of the Pauli matrices [12]. While our analysis is now much more general, our main example Θ , whose spacetime diagram is shown in figure 1a, is the classical counterpart of a CQCA.

Our paper is organized as follows: in section 1 we give our definition of a linear cellular automaton, introduce the formalism we will be working with, and state the main result: every linear cellular automaton has a fractal structure. We will also introduce the example Θ which will be the focal point of this article. In section 2, we will give an intuitive idea as to why the spacetime diagram of Θ exhibits a fractal

* We use the simple notation \mathbb{Z}_d for the cyclic group of order d , instead of $\mathbb{Z}/d\mathbb{Z}$, as we are concerned with finite groups only.




Figure 1: (a) is a projection of (b) induced by $\mathbb{Z}_4^2 \twoheadrightarrow \mathbb{Z}_2^2$.

structure. We will then proceed, in section 3, to expose an algorithm taking as input the local transition rule and outputting a description of the spacetime diagram. This allows us to compute salient features of these fractals, such as their fractal dimension and their average color.

1 DEFINITIONS

1.1 Generalities on summable automata

Monoids

We want to discuss “summable automata”, for which it makes sense to talk about the influence of a single cell on every other cell, and where the global transition function can be reconstructed by “summing” all these influences. So, if Σ denotes the alphabet, instead of the usual local transition function $\Sigma^I \rightarrow \Sigma$, a summable automaton is naturally defined by a function $\Sigma \rightarrow \Sigma^I$. What is then the minimal structure on Σ that would make such a definition work? These influences have to be “summed”, so we need an operation on Σ . Since the strip is infinite, an infinitary operation would do, but that wouldn’t give us much to work with. Instead, it seems reasonable to consider a binary operation $+$. In the same spirit, when we think of the superposition of influences coming from each cell, no notion of order between the cells is involved; even if in the one-dimensional case a natural order can be put on the cells, it would be less than clear what to do in higher dimensions. We require therefore that $+$ be associative and commutative. The last requirement comes from the fact that, given only the global transition function, we want to be able to isolate the influence of one cell; that is why we demand that $+$ have an identity element,

which makes now $(\Sigma, +)$ an Abelian monoid. Of course, in order for all of this to be relevant, the transition function has to be a morphism.

Let I be some finite subset of \mathbb{Z} and f a morphism from Σ to Σ^I . From f one can define the global transition function as an endomorphism F of $\Sigma^{\mathbb{Z}}$ by

$$F : \left(\begin{array}{ccc} \Sigma^{\mathbb{Z}} & \rightarrow & \Sigma^{\mathbb{Z}} \\ r = (r_n)_{n \in \mathbb{Z}} & \mapsto & \left(\sum_{i \in I} f(r_{n-i})_i \right)_{n \in \mathbb{Z}} \end{array} \right). \quad (1)$$

Let σ be the right shift on $\Sigma_{\mathbb{Z}}$, i.e. $\sigma(r)_n = r_{n-1}$. We have $F \circ \sigma = \sigma \circ F$, which means F is translation invariant. Also, $F(r)_n$ depends only on the values r_{n-i} for $i \in I$; since I is finite, F is a one-dimensional cellular automaton on the alphabet Σ , with neighborhood included in $-I$. Conversely, if F is an endomorphism of $\Sigma^{\mathbb{Z}}$ defining a cellular automaton over the alphabet Σ , then one can choose a neighborhood I , and define, for $i \in I$,

$$f(s)_i = F(\bar{s})_i, \quad (2)$$

where \bar{s} is the word of $\Sigma_{\mathbb{Z}}$ defined by $\bar{s}_n = \begin{cases} s & \text{if } n = 0 \\ e & \text{otherwise} \end{cases}$, e denoting the neutral element of Σ .

Groups

We will now consider the case when Σ is a (finite abelian) group. For p prime, let Σ_p be the subgroup of Σ of elements of order a power of p ; then Σ is isomorphic to $\prod_p \Sigma_p$, and every endomorphism of $\Sigma^{\mathbb{Z}}$

factorises into a product of endomorphisms of the $\Sigma_p^{\mathbb{Z}}$'s. It is therefore enough to study the case of the (abelian) p -groups: let us assume Σ is a p -group.

It is a well-known fact (see for instance section I-8 of [7]) that Σ is isomorphic to $\mathbb{Z}_{p^{k_1}} \times \mathbb{Z}_{p^{k_2}} \times \cdots \times \mathbb{Z}_{p^{k_d}}$ with $k_d \geq k_{d-1} \geq \cdots \geq k_1 = k$. Now let us consider an endomorphism α of Σ and let e_j denote $(0, \dots, 0, 1, 0, \dots, 0)$, where the 1 lies in position j . When $i \geq j$, there is a natural embedding $s_{i,j}$ of $\mathbb{Z}_{p^{k_i}}$ into $\mathbb{Z}_{p^{k_j}}$, namely the multiplication by $p^{k_j - k_i} \in \mathbb{N}$. Since e_j has order p^{k_j} , $\alpha(e_j)_i \in \mathbb{Z}_{p^{k_i}}$ has to be in the image of $s_{j,i}$ when $i \leq j$. We can therefore associate to α the endomorphism of $\mathbb{Z}_{p^k}^d$ given by the matrix $A(\alpha) \in \mathcal{M}_d(\mathbb{Z}_{p^k})$ defined by $A(\alpha)_{i,j} = p^{k_j - k_i} \alpha(e_j)_i$.

For instance, if G is $\mathbb{Z}_{32} \times \mathbb{Z}_4 \times \mathbb{Z}_2$, and α is defined by $\alpha(1,0,0) = (3,2,1)$, $\alpha(0,1,0) = (24,0,1)$ and $\alpha(0,0,1) = (16,2,0)$, then the corresponding matrix of $\mathcal{M}_3(\mathbb{Z}_{32})$ would be

$$A(\alpha) = \begin{pmatrix} 3 & 3 & 1 \\ 16 & 0 & 1 \\ 16 & 2 & 0 \end{pmatrix}.$$

Let us give a summary of the construction we have just exposed.

Proposition 1. *For every finite abelian p -group G and endomorphism α of G , there are positive integers k and d , an embedding s of G into $\mathbb{Z}_{p^k}^d$, and an endomorphism $A(\alpha)$ of $\mathbb{Z}_{p^k}^d$ such that the following diagram commutes:*

$$\begin{array}{ccc} G & \xrightarrow{\alpha} & G \\ \downarrow s & & \downarrow s \\ \mathbb{Z}_{p^k}^d & \xrightarrow{A(\alpha)} & \mathbb{Z}_{p^k}^d \end{array} \quad (3)$$

This implies that to study the behaviour of CA on abelian groups, it is enough to study the case where these groups are of the form $\mathbb{Z}_{p^k}^d$.

R-modules

We will actually consider the more general case where R is a finite commutative ring, and Σ is a free R -module of dimension d , i.e. isomorphic to R^d . The first reason for doing so is that it does not complicate the mathematics. It will also appear more efficient to understand, for instance, \mathbb{F}_{2^4} as a 1-dimensional vector space over itself than as a 4-dimensional vector space over \mathbb{F}_2 : the former simply bears more information, and therefore implies more restrictions on the form of a CA, so that more can be deduced.

For any ring B , $B[u, u^{-1}]$ denotes the ring of Laurent polynomials over B ; it is the ring of linear combinations of integer powers (negative as well as nonnegative) of the unknown u . Applying this to $B = \text{Hom}_R(\Sigma)$, we can associate to the function f the Laurent polynomial $\tau(f) \in \text{Hom}_R(\Sigma)[u, u^{-1}]$ defined by

$$\tau(f) = \sum_{n \in \mathbb{Z}} f(\cdot)_n u^n. \quad (4)$$

τ is an isomorphism of R -algebras between the linear cellular automata on the alphabet Σ with internal composition rules $(+, \circ)$ and

$\text{Hom}_R(\Sigma)[u, u^{-1}]$, which can be identified with $\mathcal{M}_d(R[u, u^{-1}])$ because $\Sigma \simeq R^d$; we are going to think and work in this former algebra, so from now on a linear cellular automaton $T = \tau(f)$ will be for us an element of $\mathcal{M}_d(R[u, u^{-1}])$.

1.2 Related work

Many papers have been published about the fractal structure of cellular automata spacetime diagrams. We give here a short review and point out the differences to our approach. When we mention d and k we are referring to $\mathcal{M}_d(\mathbb{Z}_k[u, u^{-1}])$.

[15] In this work, Willson considers the case $d = 1, k = 2$. In order to determine the fractal dimension of the spacetime diagram, he analyses how blocks of length n in the configuration of time step t are mapped to such blocks in steps $2t$ and $2t + 1$, a technique we also use in section 3.

[13, 14] Takahashi generalizes Willson's work to the case $d = 1$, with no restriction on the value of k .

[5, 6] Haeseler, Peitgen and Skordev study the fractal time evolution of CAs with special scaling properties, the weakest of them being "weakly p -Fermat", where p is some integer, which includes the case $d = 1, k = p$. Let us briefly introduce the p -Fermat property and show why the CAs that we study do not have to be p -Fermat. Let π_p be the scaling map

$$\pi_p(\xi)_x = \begin{cases} \xi_y & \text{if } x = py \\ e & \text{otherwise} \end{cases} . \quad (5)$$

A CA T is weakly p -Fermat if for all $s \in \Sigma, n \in \mathbb{N}$ and $x \in \mathbb{Z}$, $T^{np}(\bar{s})_x = e \Leftrightarrow \pi_p T^n(\bar{s})_x = e$.

Let us now consider

$$\Theta = \begin{pmatrix} 0 & 1 \\ 1 & u^{-1} + 1 + u \end{pmatrix} \in \mathcal{M}_2(\mathbb{Z}_2[u, u^{-1}]) . \quad (6)$$

We will use this example throughout the paper. It generates the time evolution depicted in Figure 2a. A general nearest-neighbor p -Fermat CA produces a time evolution that reproduces itself after p steps in at most three copies located at positions $\{-p; 0; p\}$. After $2p$ steps we have five copies at most. This creates areas

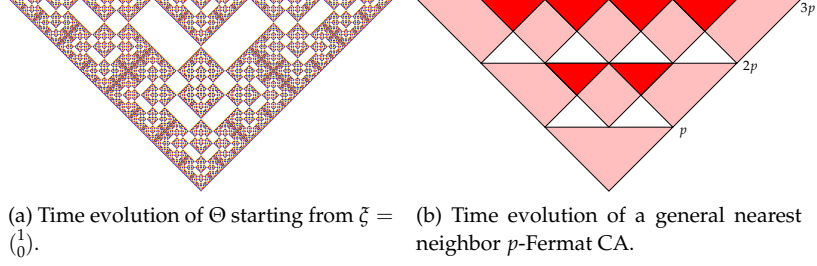


Figure 2: This figure shows that Θ cannot be a p -Fermat CA. In a p -Fermat CA at least the white areas are filled by the neutral element e ; Θ has a different pattern.

filled with the neutral element e shared by all p -Fermat CAs for a fixed p . In figures 2a and 2b we can easily see that Θ does not exhibit these areas; therefore it is not p -Fermat. Furthermore p -Fermat CAs that are not periodic are irreversible, while we also allow reversible CAs, Θ being again one example.

[2, 1] Allouche, Haeseler, Peitgen and Skordev study recurrences in the spacetime diagram of linear cellular automata, from the angle of k -automatic sequences, which we will not define in this paper. However they require Σ to be an Abelian ring and the CA to be a ring homomorphism, which is again essentially the case $d = 1$.

[10, 11] Moore studies CAs with an alphabet A on a staggered space-time, where every cell c is only influenced by two cells a and b of the last time step. The update rule is $c = a \bullet b$. He requires (A, \bullet) to be a quasigroup and studies different special cases. First let us note that these CAs are either irreversible or trivial, while ours don't have to be. Thus, although it is possible to bring our CAs in the form of staggered CAs, the results of Moore do not apply. In his setting, our CAs would be of the form $c = a \bullet b = f(a) + g(b)$ for some homomorphisms f and g . For (A, \bullet) to be a quasigroup means

$$\forall a, b \in A \exists! x, y \in A \quad a \bullet x = b \wedge y \bullet a = b. \quad (7)$$

In our case, these equalities translate respectively as $g(x) = b - f(a)$ and $f(y) = b - g(a)$. The right-hand sides can be arbitrary elements of A , therefore f and g have to be isomorphisms, as indeed required in [10].

The angle of study of Moore is also different: he does not exactly study the fractal properties of the spacetime diagram, but rather the complexity of the prediction — “What will be the state of this cell after t steps?”. Describing the spacetime diagram with a matrix substitution system is one way of proving that prediction is an easy task — for instance it makes it NC.

- [8] Macfarlane uses Willson’s approach and generalizes parts of it to some examples of matrix-valued CAs, including Θ . However, the transition matrix is obtained heuristically — “by scrutiny of figure 9” — from the spacetime diagram, instead of being algorithmically derived from the transition rule (as in the present work). The conclusion (section 6) suggests that the analysis of Θ is easily generalizable to matrices of various sizes over various rings, so in a sense the present article is but an elaboration of the concluding remark of [8], although we have to say we do not find this generalization to be that obvious.

The heart of our proof is in section 3. In a nutshell, whereas most of the techniques used in our article can be traced back to older articles, the new one that allows us to extend the analysis to a larger class of automata is the introduction of α in Equation (19). The idea in doing so is to get rid of the complicated noncommutative ring structure and go back to a simple linear recurrence, as stated in Proposition 4. Since a linear recurrence is precisely where the analysis started from, it could seem at first sight that nothing is gained in the process, but the new recurrence actually does not define a cellular automaton. Instead of defining line $n + 1$ from line n , it cuts right through to line mn , thus establishing a scaling property.

1.3 Different CAs

While we use Θ , which has very special properties (being reversible, over a field of characteristic two, and described by a 2×2 matrix), as our example throughout the paper, the analysis applies of course to all other linear CAs. In this section we give a short overview over the variety of spacetime diagrams these CAs generate. Let us start with

small changes to Θ . For our first example we keep $m = k = 2$, but change the determinant to u . We only change one entry of the matrix:

$$\Theta_u = \begin{pmatrix} 0 & u \\ 1 & u^{-1} + 1 + u \end{pmatrix}. \quad (8)$$

The spacetime diagram is displayed in figure 3a and shows how much difference a small change in the update rule can make for the spacetime diagram.




Figure 3: Spacetime diagrams of non-clifford CAs.

Let us now modify Θ in a more subtle fashion:

$$\Theta_{k=4} = \begin{pmatrix} 0 & 1 \\ 1 & u^{-1} + 1 + u \end{pmatrix}. \quad (9)$$

The hidden difference with Θ is the underlying ring, which has now been extended from \mathbb{Z}_2 to \mathbb{Z}_4 . $\Theta_{k=4}$ contains in some sense more information than Θ , since Θ is induced from $\Theta_{k=4}$ by the projection $\mathbb{Z}_4 \rightarrow \mathbb{Z}_2$. Consequently, the spacetime diagram of Θ is nothing but a projection of that of $\Theta_{k=4}$, as illustrated in figure 1b.

The last CA we want to present lies in $\mathcal{M}_2(\mathbb{F}_4[u, u^{-1}])$, where \mathbb{F}_4 is the finite field of order 4, here identified with $\mathbb{F}_2[\omega]/(\omega^2 + \omega + 1)$. The corresponding matrix is

$$T_{\mathbb{F}_4} = \begin{pmatrix} 0 & \omega \\ u^{-1} & (\omega + 1)u^{-1} + \omega + u \end{pmatrix}. \quad (10)$$

If one wants to avoid calculations in \mathbb{F}_4 , this CA can be translated to a

CA in $\mathcal{M}_4(\mathbb{Z}_2[u, u^{-1}])$, namely

$$\tilde{T}_{\mathbb{F}_4} = \begin{pmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 1 \\ u^{-1} & 0 & u^{-1} + u & u^{-1} + 1 \\ 0 & u^{-1} & u^{-1} + 1 & 1 + u \end{pmatrix}.$$

Its spacetime diagram, as can be seen shown in figure 3b, contains patches of checkerboard pattern. Somehow, they trivialise most of the usual properties of the figure: for instance they make its fractal dimension 2, even if the fractal structure can hardly be considered trivial. In order to access more interesting properties, it is possible to blank this pattern out, considering it as just “another shade of white”. This can be trivially done on the matrix substitution system, by removing the states from which the blank state is inaccessible.

1.4 Coloured spacetime diagrams

The mainstream setting when studying the fractal structure of spacetime diagrams is monochromatic; we introduce colors in the picture.

Instead of considering simple compact subsets of the plane, we will have a finite set of colors \mathcal{C} and compact subsets of $(\mathbb{R}^2)^{\mathcal{C}}$. Let $b \notin \mathcal{C}$ be the additional “blank” color and $c : \Sigma \rightarrow \mathcal{C} \cup \{b\}$ a coloring of Σ such that $c(0) = b$. To determine a *colored spacetime diagram*, we need furthermore to be given an automaton $T \in \mathcal{M}_d(R[u, u^{-1}])$, an initial state $\xi \in R^d$, and an integer n . The corresponding colored spacetime diagram is then the rescaled diagram obtained by iteratively applying T n times on ξ .

Formally, for $n, i, j \in \mathbb{N}$, let $S_{n,i,j}$ be the full square centred in $\frac{1}{n}(i, j)$ and whose edges, parallel to the axes, are of length $\frac{1}{n}$. To each positive integer n and color $c \in \mathcal{C}$ is associated a compact subset of the plane $\mathcal{P}_n(c)$ which is the union of the $S_{n,i,j}$'s such that $0 \leq j \leq n$ and $c(T^j(\xi)_i) = c$. The colored spacetime diagram of order n is then the function $\mathcal{P}_n : c \mapsto \mathcal{P}_n(c)$. A sequence of colored patterns $(\mathcal{P}_n)_{n \in \mathbb{N}}$ of spacetime diagrams is said to converge to some colored pattern \mathcal{P}_∞ if for every $c \in \mathcal{C}$, $(\mathcal{P}_n(c))_{n \in \mathbb{N}}$ converges to $\mathcal{P}_\infty(c)$ for the Hausdorff distance.

We can now state our main result.

Theorem 1. *Let G be a finite abelian p -group. For every cellular automaton over G that is also a group homomorphism, there exists a positive integer m*

such that for every fixed initial state the coloured spacetime diagrams of order p^{mn} converge when n goes to infinity.

In general, to know about the fractal structure of a cellular automaton over some finite group G , write G as a product of p -groups and study each p -component of the spacetime diagram independently; according to Theorem 1, each component generates a fractal pattern. Then, since the logarithms of the prime numbers are rationally independent, it is possible to find a sequence of resized spacetime diagrams that converges towards a superposition of these different components with arbitrary independent rescaling coefficients, but there is no direct generalization of the theorem. For instance, even in the simple case of Pascal's triangle modulo 6, there is no real number $\alpha > 0$ such that the diagrams of order $\lfloor \alpha^n \rfloor$ converge; however those of order t_n will converge as soon as the fractional parts of $\log_3(t_n)$ and $\log_2(t_n)$ both converge, and then their limits determine the limit pattern. The situation is very briefly described in section 5 of [14].

Matrix substitution systems

We will show how to find a suitable description of the limit pattern in the rest of this article. We now explain exactly what it means to generate a colored picture by rules of substitution, and how to take the limit of all these pictures. This is a generalization of the usual monochromatic description that can be found for instance in [9, 15, 5], and which corresponds in our setting to the case where all the colors are mapped to "black".

Let V be a finite alphabet; because we want colors, compared with the usual definition of a matrix substitution system, we don't have to include a special "empty" letter. A matrix substitution system is then a function $\mathcal{D} : V \rightarrow V^{\llbracket 1;r \rrbracket^2}$; for some integer r . Together with a set of colors \mathcal{C} and a coloring $\mathfrak{c} : V \rightarrow \mathcal{C}$, it defines colored patterns, much in the same way cellular automata do. With the previous notations, at each step n , the pattern \mathcal{P}_n is the union of squares $S_{r^n, i, j}$ of different colors, for different i 's and j 's; each one of them is indexed by some letter in V .

Then at step $n + 1$, each colored square of color c indexed by $v \in V$ present in the n th step pattern is replaced by r^2 smaller squares that pave it; these smaller squares are given by $\mathcal{D}(v)$ and indexed accordingly. To such a matrix substitution system we can associate a multi-

graph $\Gamma = (V, E)$ where the set of vertices is V and we put as many edges from v to w as there are w 's in $\mathcal{D}(v)$.

A *plain* matrix substitution system is one of the usual kind: no coloring, and V contains a special letter ε such that $\mathcal{D}(\varepsilon)$ is a matrix full of ε 's and $\mathfrak{c}(\varepsilon) = b$. In the multigraph associated to a plain matrix substitution system, ε is excluded from the set of vertices.

We want to generalize the usually property of convergence of the patterns defined by plain matrix substitution systems. This will be done by the conjunction of the two following propositions. Let us first remind some notions on graphs: the *period* of a graph is the greatest common divisor of the lengths of all the cycles in Γ ; a graph is *aperiodic* if it has period 1.

Proposition 2. *If every strongly connected component of Γ is aperiodic, then $(\mathcal{P}_n)_{n \in \mathbb{N}}$ converges.*

Proof. To each color $c \in \mathcal{C}$ we associate the plain matrix substitution system \mathcal{D}^c , obtained from \mathcal{D} simply by turning some letters into ε . For $v \in V$, let $X_c(v)$ be the set of integers n such that there exists a path of length n in Γ connecting v to a letter of the color c . Since the strongly connected component containing v is aperiodic, $X_c(v)$ is either finite or cofinite. Those letters $v \in V$ such that $X_c(v)$ is finite are sent to ε , and this defines \mathcal{D}^c . If $X_c(v)$ is finite and v' can be reached from v , then $X_c(v')$ is also finite; therefore, \mathcal{D}^c is indeed a substitution system. Let M be such that for every $v \in V$, either $X_c(v)$ or its complement is strictly bounded by M .

Let us now compare two sequences of figures. The first one is $(\mathcal{P}_n(c))$, the subpattern of color c defined by \mathcal{D} . The second one is (\mathcal{P}_n^c) , the one obtained from \mathcal{D}^c ; we know that it converges to some compact \mathcal{D}_∞^c . By construction, $\mathcal{P}_{n+M}(c)$ is included in \mathcal{P}_n^c , and for every black square of \mathcal{P}_n^c , there is a black subsquare in $\mathcal{P}_{n+M}(c)$. The Hausdorff distance between $\mathcal{P}_{n+M}(c)$ and \mathcal{P}_n^c therefore converges to 0, so that $(\mathcal{P}_n(c))$ converges to \mathcal{D}_∞^c . \square

For a graph Γ , let $\Gamma^k = (V, E^k)$ where E^k is the set of couples (v, w) such that there exists in Γ a path of length k from v to w .

Proposition 3. *For every (multi)graph Γ , there exists k such that every strongly connected component of Γ^k is aperiodic.*

Proof. Each strongly connected component Δ of Γ has a period $p(\Delta)$, so that $\Delta^{p(\Delta)}$ is aperiodic. Let k_0 be the least common divisor of the

$p(\Delta)$'s; then each strongly connected component of Γ induces an aperiodic graph in Γ^{k_0} , but it is possible that, in the process, it broke down into several connected components, so that Γ^{k_0} might not have the required property. The procedure then has to be repeated from Γ^{k_0} to obtain $\Gamma^{k_0 k_1}$, and so on. Since the strongly connected components of $\Gamma^{k_0 \cdots k_{i+1}}$ are included in those of $\Gamma^{k_0 \cdots k_i}$, this process reaches a fixed point, which is a graph with the required property. \square

Ergo, a colored matrix substitution system defines a convergent colored pattern when considering the steps that are a multiple of some well-chosen integer m . So, in order to prove Theorem 1, all we need to do is find such a substitution system. This will be done in a special case in the next section, and in the general case in section 3.

2 A SPECIAL RECURSION SCHEME FOR Θ

The aim of this section is to give the most direct and natural explanation of the fractal structure generated by Θ that we are aware of. Modulo some caveat, it applies effortlessly to all invertible elements T of $\mathcal{M}_2(R[u, u^{-1}])$, where R is a finite Abelian ring of characteristic 2. This section is not vital to the proof of the general case presented in 3, and can therefore be skipped by the impatient reader.

We will deduce informally the basic structure of the spacetime diagrams from a simple recursion relation for the 2^n th powers of T . The characteristic polynomial of T , $P_T(X)$, is equal to $X^2 + (\text{tr } T)X + \det T$. According to Cayley-Hamilton theorem, $P_T(T) = 0$, so $T^2 + (\text{tr } T)T + (\det T)\mathbb{I} = 0$. Multiplying this equation by T^{-1} , we get $T = (\det T)T^{-1} + (\text{tr } T)\mathbb{I}$. Let us denote $\tilde{T} = (\det T)T^{-1}$, which we will name the *dual* of T ; since we are in characteristic 2, by repeatedly taking the square of this equality, we obtain

$$\forall n \in \mathbb{N} \quad T^{2^n} = \tilde{T}^{2^n} + (\text{tr } T)^{2^n} \mathbb{I}. \quad (11)$$

Taking the trace of this equation, we get $\text{tr } T^{2^n} = \text{tr } \tilde{T}^{2^n}$; in particular, $\text{tr } T = \text{tr } \tilde{T}$ so Equation (11) is also valid when swapping T and \tilde{T} . Let I_T be a finite set, and the λ_i 's elements of R such that $\text{tr } T = \sum_{i \in I_T} \lambda_i u^i$. Then we have

$$\forall n \in \mathbb{N} \quad (\text{tr } T)^{2^n} = \sum_{i \in I_T} (\lambda_i)^{2^n} u^{2^n i}. \quad (12)$$

We do not yet specify the initial state; as a matter of fact, it will prove to be largely irrelevant. The only thing we ask for now is that it is nontrivial (and finite).

Consider for instance Θ ; we have $\det \Theta = 1$ and $\lambda_i = \chi_{\{-1,0,1\}}(i)$. We start with the spacetime diagram corresponding to 2^n steps; it is rescaled to a triangle with vertex coordinates $\{(0,0), (-1,1), (1,1)\}$. Taking equations (11) and (12), we can see that the state at the 2^n th time step can be decomposed into a sum of several copies of the initial state (the positions are governed by the coefficients of the trace) and a configuration that can be derived by applying \tilde{T}^{2^n} to the initial state. In the next 2^n steps, this configuration will contract itself to the initial state, which is shifted according to $\log_u \det T$, as \tilde{T} is the inverse of T composed with the shift $(\det T)\mathbb{I}$. The copies of the original initial state evolve according to T . This is illustrated in Figure 4. The figure




Figure 4: The whole figure is the sum of $|I| + 2$ parts.

suggests to divide the spacetime diagram into four parts A , B , C , and D as shown in Figure 5a which overlap only on a single cell strip at the borders.

A_2 , A_3 , and A_4 are copies of A_1 , and V is marked by an upside down A because it is the reverse evolution of the initial state under the CA $\tilde{\Theta}$ — so actually it should logically be named \tilde{V} or \check{V} , but since it always appears upside-down while A always appears straight on its




Figure 5: The first and the second substitution rules.

feet, there is no risk of confusion.

Let us assume that the sequence of rescaled spacetime diagrams up to step 2^n actually converges. Then that means A_1 should be, in the limit, a copy of the whole picture A , downsized by a factor 2, so we rename it A . This gives us the first substitution rule, represented in Figure 5a. The other three parts are still unknown, and we will name these patterns B , C and D . Equation 11 tells us what the other substitution rules are.

Since Equation 11 remains true after swapping T and \tilde{T} , V admits likewise a partition into V , \mathcal{A} , \mathcal{C} and \mathcal{D} . Summing all the parts shown in Figure 4, we get the top rightmost pattern of Figure 9. Superimpo-




Figure 6: Second step of the decomposition.

sing our first substitution rule (Figure 5a) with our second step of the decomposition (Figure 6), we get the new substitution rules represented by Figures 5b, 7a and 7b.

All the other substitution rules are deducible from these ones. First they are linear: for instance the substitutions for $C + D$ is the sum of the substitutions for C and D , as shown in Figure 8a. We don't know *a priori* what the sum of two patterns is, but we know that summing a pattern with itself should give 0, for we are working in characteristic




Figure 7: The third and the fourth substitution rules.

2. In this case $A + A$ and $B + B$ cancel out. In figure 9, one can see how the characteristic white spaces emerge from the above substitution rule.




Figure 8: The fifth and sixth substitution rules are the superpositions of the third and fourth and the first and second rules respectively. One can see, how white spaces emerge because even superpositions of the same pattern cancel out.

The problem with this scheme is that what is happening goes beyond simple juxtaposition of patterns. There can be cancellation at the border between patterns. And, sure, we know $C + C$ is blank, but how do we know $C + D$, for instance, is not? Well, generally we don't. If the initial state is itself blank, then the whole figure would be, and all tiles being blank is certainly a fixed point for all the substitution rules.

In this case, however, everything turns out well. It should first be noticed that in every part of the picture not tagged as “blank”, an A pattern can be found by refining a few more steps. Formally, let $G = (V, E)$ be the graph whose vertices are the different possible tiles (i.e., in this case, $A, B, C, D, V, g, \emptyset, \emptyset$, and the sums modulo 2 of tiles having compatible shapes, including the blank tile 0), and edges represent the transition rule in the following way: each vertex has four edges coming out of it, each one pointing to one of its subtiles. In our case, the graph has the property that the set of vertices accessible from A , minus 0, form a strongly connected component.

We may then distinguish two cases: either A has a point in its interior, or it has points only on its border triangle. In the first case, the unique non-empty compact defined by the substitution rule is actually the figure we're looking for. Indeed, it follows from the property of connexity — cf. Proposition 2 — that every non-zero tile actually appearing in the decomposition of the figure has a non-empty interior. Thus, no matter what happens at the boundary between tiles, the figure constructed this way will always converge to the same compact.




Figure 9: Three decomposition steps for the spacetime diagram of Θ . One can clearly see the characteristic white spaces emerging.

3 RECURSION AND MATRIX SUBSTITUTION SYSTEM

We will now present a general method to calculate the fractal dimension and average colour of the spacetime diagrams of linear CAs in $\mathcal{M}_d(\mathbb{Z}_{p^l}[u, u^{-1}])$. We will again illustrate the method with our example Θ , whereas the derivation is carried out for the general case. Thus the algorithm works as well on all CAs obeying our definition, e.g. the CAs presented in section 1.3, as it works on Θ . Of course with larger neighborhoods and groups of higher order the substitution system becomes larger and larger, so that one might want to use a computer to derive the substitution system.

Our approach is the following: from the minimal polynomial Π of the CA T (or any other polynomial fulfilling $\Pi(T) = 0$) we derive a recursion relation for the T_x^y 's, the coefficients in u^x of T^y . We then forget about every other piece of information we might have on T , to concentrate only on this recursion: this shows that the fractal structure, except for contingent blank spaces, can be essentially derived just from the minimal polynomial of T . We further develop our recursion scheme for T until we can express every T_x^y in terms of the T_i^j of the first m time steps with coefficients $\alpha_j(x - i, y)$. With a simple grouping of cells we deduce a matrix substitution system that enables us to generate the spacetime diagram of step $t = k^{n+1}$ directly from step $t = k^n$. Using this substitution system we can calculate the fractal dimension,

the average colouring, and given an initial state also the whole space time diagram.

Now let $\Pi(X) \in R[u, u^{-1}][X]$ be a monic polynomial such that $\Pi(T) = 0$:

$$\Pi(X) = X^m - \sum_{j=0}^{m-1} \lambda_{\Pi,j} X^j. \quad (13)$$

According to the Cayley-Hamilton theorem, which we can apply in our case because T is an endomorphism of a finite-dimensional free module over an abelian ring (see Theorem 3.1. of [7]), the characteristic polynomial of T fulfills this condition, therefore we can always find such a polynomial. Let \mathcal{I} be the finite set of exponents i 's such that the coefficient in u^i of Π , seen as an element of $R[X][u, u^{-1}]$, is nonzero, so that we can write $\lambda_{\Pi,j} = \sum_{i \in \mathcal{I}} \lambda_{\Pi,i,j} u^i$. \mathcal{I} is not to be confused with the neighbourhood of the CA, I , which won't play any role from now on.

For any $x, y, n \in \mathbb{N}$, $(x + y)^p \equiv x^p + y^p[p]$, and if $x \equiv y[p^n]$ then $x^p \equiv y^p[p^{n+1}]$. Therefore, for any $x, y \in R$ and $n \in \mathbb{N}$, $(x + y)^{p^{n+l-1}} = (x^{p^n} + y^{p^n})^{p^{l-1}}$. Since the powers of T commute pairwise, we get

$$\begin{aligned} T^{p^{n+2(l-1)}m} &= \left(\sum_{j=0}^{m-1} \lambda_{\Pi,j}^{p^{n+l-1}} T^{p^{n+l-1}j} \right)^{p^{l-1}} \\ &= \left(\sum_{j=0}^{m-1} \left(\sum_{i \in \mathcal{I}} \lambda_{\Pi,i,j}^{p^n} u^{p^n i} \right)^{p^{l-1}} T^{p^{n+l-1}j} \right)^{p^{l-1}}. \end{aligned} \quad (14)$$

For each i, j , the sequence $(\lambda_{\Pi,i,j}^{p^n})$ is ultimately periodic. There exist therefore integers N and M such that for all i, j , $\lambda_{\Pi,i,j}^{p^{M+N}} = \lambda_{\Pi,i,j}^{p^M}$. Let $k = p^N$; substituting n by $M + Nn$ in (14), we get

$$T^{k^n p^{M+2(l-1)}m} = \left(\sum_{j=0}^{m-1} \left(\sum_{i \in \mathcal{I}} \lambda_{\Pi,i,j}^{p^M} u^{k^n p^M i} \right)^{p^{l-1}} T^{k^n p^{M+l-1}j} \right)^{p^{l-1}}. \quad (15)$$

Hence, if we note $m' = p^{M+2(l-1)}m$ and expand this equation, we find that there is some finite subset \mathcal{I}' of \mathbb{Z} and some elements $\mu_{i,j}$ of R , for $i \in \mathcal{I}'$ and $j \in \llbracket 0; m' - 1 \rrbracket$, such that for all $n \in \mathbb{N}$,

$$T^{k^n m'} = \sum_{j=0}^{m'-1} \sum_{i \in \mathcal{I}'} \mu_{i,j} u^{k^n i} T^{k^n j}. \quad (16)$$

We have now used everything we needed to know from the multiplicative structure on the ring of matrices. As announced at the end of section 1.2, we will now get rid of it and concentrate only on the linear recurrence relation that we have just derived. Remember that $T^j \in \mathcal{M}_d(R[u, u^{-1}])$, and we are interested in the coefficient of T^j in u_i , denoted T_i^j , so that $T^j = \sum_{i \in \mathcal{I}} T_i^j u^i$. We thus get the following relation: $T_x^{k^n m' + y} = \sum_{i \in \mathcal{I}'} \sum_{j=0}^{m'-1} \mu_{i,j} T_{x-k^n i}^{y+k^n j}$, which we rewrite in this form:

$$T_x^y = \sum_{(i,j) \in \mathcal{I}' \times \llbracket 0; m'-1 \rrbracket} \mu_{i,j} T_{x+f_{i,j}(y)}^{g_{i,j}(y)} \quad (17)$$

where $f_{i,j}(y) = -k^n i$ and $g_{i,j}(y) = y - k^n(m' - j)$, which of course works with any n , but we will choose $n = \lfloor \log_k \frac{y}{m'} \rfloor$. In order to emphasise that the rest of the proof will use only a minimal structure, we state in the next proposition what will actually be proven, and change the notation from T , which was an element of $\mathcal{M}_d(R[u, u^{-1}])$, to Ξ , an element of a more arbitrary R -module. It is straightforward to check that T fulfills the hypotheses of the proposition.

Proposition 4. *Let M be a finite R -module, k a positive integer, Λ a finite set of indices, and for $i \in \Lambda$, $\mu_i \in R$, $f_i : \llbracket m; +\infty \rrbracket \rightarrow \mathbb{Z}$ and $g_i : \llbracket m; +\infty \rrbracket \rightarrow \mathbb{N}$ such that for all $y \in \llbracket m; +\infty \rrbracket$ and $t \in \llbracket 0; k-1 \rrbracket$,*

- $g_i(y) < y$;
- $f_i(ky + t) = kf_i(y)$ and $g_i(ky + t) = kg_i(y) + t$.

For $x \in \mathbb{Z} \times \mathbb{N}$, let $\Xi_x^y \in M$ be such that when $y \geq m$,

$$\Xi_x^y = \sum_{i \in \Lambda} \mu_i \Xi_{x+f_i(y)}^{g_i(y)}. \quad (18)$$

Then there exists a finite set E and a function $e : \mathbb{Z} \times \mathbb{N} \rightarrow E$ such that

- Ξ_x^y is a function of $e(x, y)$;
- for $s, t \in \llbracket 0; k-1 \rrbracket$, $e(kx + s, ky + t)$ is a function of s, t , and $e(x, y)$.

The introduction of a new function e in this proposition comes from the need of a scaling property, expressing that the state at point $(kx + s, ky + t)$ can be deduced from the state at point (x, y) . Such a property does not follow immediately from Equation (18), but it is possible

to expand the state space from M to E , and to put more information into e than into Ξ , so as to fulfill the scaling property. An immediate consequence of this proposition is that the spacetime diagrams of Ξ_x^y of order k^n can be described by colored matrix substitution systems, so that Theorem 1 will follow from Proposition 3. Let us now prove Proposition 4.




Figure 10: Recursive calls to (18) for $T = \Theta$.

If $y \geq m$, we can recursively apply Equation (18) to give an expression of Ξ_x^y in terms of a linear combination of $\Xi_{x'}^{y'}$'s with $y' < y$.

Let us consider the case $T = \Theta$ as an example. In this case, $R = \mathbb{F}_2$, $k = m = 2$, and $\Pi_\Theta(X) = X^2 + (u^{-1} + 1 + u)X + 1$. Equation (18) gives $\Xi_1^3 = \Xi_1^1 + \Xi_0^2 + \Xi_1^2 + \Xi_2^2$. Whereas the decomposition halts here for Ξ_1^1 , it goes on for the other three terms of the sum, again with $n = 0$. For instance, Ξ_0^2 is transformed into $\Xi_0^0 + \Xi_{-1}^1 + \Xi_0^1 + \Xi_1^1$. This gives the tree of recursive calls in Figure 10, where each node is labelled (x, y) for Ξ_x^y ; in that particular case, because $R = \mathbb{F}_2$, the only possible coefficients are 0 and 1, so they are simply represented by the absence or the presence of the corresponding term.

Starting from any point $(x, y) \in \mathbb{Z} \times \mathbb{N}$ and using recursively equation (18), we get to the expression

$$\Xi_x^y = \sum_i \sum_{j=0}^{m-1} \alpha_{\Pi, i, j}(x, y) \Xi_i^j, \quad (19)$$

which we take as a definition of $\alpha_{\Pi, i, j}(x, y)$. On Figure 10, this corresponds to (parity) counting the leaves tagged (i, j) ; for instance, there are 4 leaves tagged $(1, 1)$, so $\alpha_{\Pi_\Theta, 1, 1}(1, 3) = 0$ and $(1, 1)$ does not appear as a node in Figure 10.

Since the relation (18) is invariant under translations of the parameter x , we have $\alpha_{i, j}(x, y) = \alpha_{0, j}(x - i, y)$. Noting $\alpha_j := \alpha_{0, j}$, we get the




Figure 11: Ξ_1^3 as a sum of the simplest terms, when $T = \Theta$.

following equation:

$$\Xi_x^y = \sum_{i \in \mathbb{Z}} \sum_{j=0}^{m-1} \alpha_j(x-i, y) \Xi_i^j. \quad (20)$$

Let us now show that $\alpha \cdot (kx+s, ky+t)$, for $s, t \in \llbracket 0; k-1 \rrbracket$, is a function of s, t , and the $\alpha \cdot (x', y)$'s, where x' ranges over some neighbourhood of x . By substituting x with $kx+s$ and y with $ky+t$ in Equation (18), we get

$$\Xi_{kx+s}^{ky+t} = \sum_{i \in \Lambda} \mu_i \Xi_{k(x+f_i(y))+s}^{kg_i(y)+t}. \quad (21)$$

Because the indices and exponents of Ξ on the left and right side of this equation have undergone the same transformation $(x, y) \mapsto (kx+s, ky+t)$, we arrive recursively at this point

$$\Xi_{kx+s}^{ky+t} = \sum_{i \in \mathbb{Z}} \sum_{j=0}^{m-1} \alpha_j(x-i, y) \Xi_{ki+s}^{kj+t} \quad (22)$$

which we want to compare to the following equation, directly deduced from (20):

$$\Xi_{kx+s}^{ky+t} = \sum_{i \in \mathbb{Z}} \sum_{j=0}^{m-1} \alpha_j(kx+s-i, ky+t) \Xi_i^j \quad (23)$$

Of course, there can be terms in (22) with $kj+t \geq m$, so that the decomposition is not over: it then needs to be performed to its end. For instance, in our example $T = \Theta$, if instead of the tree of recursive calls for Ξ_1^3 we want that of $\Xi_3^6 = \Xi_{2 \times 1 + 1}^{2 \times 3}$, we know that the upper part of the tree will be Figure 12.

As for the lower part, the decomposition is not over, so it needs to be performed to its end. Thanks to linearity, we don't need to remember the whole tree, just its leaves. As is well known, a tree grows from its




Figure 12: Recursive calls to (18) for $T = \Theta$, minus some leaves.

leaves, and in the same way we need only apply the transformation $(i, j) \mapsto (2i + 1, 2j)$ to the tree in Figure 11 in order to find the correct decomposition for Ξ_3^6 , which gives Figure 13.




Figure 13: Trees grow from their leaves.

However, even keeping track only of the leaves would require an unbounded alphabet, that is why we defined $\alpha_{\Pi,j}(x, y) = \alpha_{\Pi,i,j}(x + i, y)$, i.e. from the tree for Ξ_x^y we want to keep track only of the leaves labelled by $(0, j)$. Alas, that will not work so easily: in Figure 13, $(0, 1)$ appears once as a son of $(-1, 2)$, which comes from node $(-1, 1)$ in Figure 11. So it appears we cannot just forget about everything else.

As we just saw on one example, $\alpha_{\Pi,\cdot}(kx + s, ky + t)$ typically does not depend on $\alpha_{\Pi,\cdot}(x, y)$. Instead the final decomposition of (22) will relate the coefficients $\alpha_{\Pi,j}(kx + s - i, ky + t)$ of the Ξ_i^j in (23) to sums of the $\alpha_{\Pi,j}^k(x - i, y)$'s. Therefore $\alpha_{\Pi,\cdot}(kx + s, ky + t)$ depends on the $\alpha_{\Pi,\cdot}(x + i, y)$'s for i ranging over some finite set \mathcal{S} depending only on Π . However, this is not much of a problem, as a simple grouping will take care of it — a technique commonly attributed to [15]. Let us show that \mathcal{S} is finite. Since $j \in \{0; \dots; m - 1\}$ and $t \in \{0; \dots; k - 1\}$, the $kj + t$ appearing as an exponent of Ξ in (22) is in $\{0; \dots; km - 1\}$. We therefore have to use at most $(k - 1)m$ recursive calls to (18) in order to get down to coefficients Ξ_x^y with $y < m$, each one of them decreasing the exponent by at least 1. Each one of them also increases the index by

$f_i(y)$; since both Λ and $\llbracket 0; km - 1 \rrbracket$ are finite, the set of possible $f_i(y)$'s is also bounded by some M , and the total variation in the index, i.e. \mathcal{S} , is then bounded by $(k - 1) mM$; let us say $\mathcal{S} \subseteq \llbracket d_{\min}; d_{\max} \rrbracket$.

Let us now introduce $\beta_{\Pi, \cdot}(x, y) = (\alpha_{\cdot}(x - i, y))_{i \in \mathcal{S}'}$, where $\mathcal{S}' = \llbracket \delta_{\min}; \delta_{\max} \rrbracket$, δ_{\max} being such that $\delta_{\max} \geq d_{\max} + \left\lceil \frac{\delta_{\max}}{k} \right\rceil$ and δ_{\min} such that $\delta_{\min} \leq d_{\min} - 1 + \left\lceil \frac{\delta_{\min} + 1}{k} \right\rceil$. This time, for $s, t \in \{0; \dots; k - 1\}$, $\beta_{\Pi, \cdot}(kx + s, ky + t)$ does really depend only on $\beta_{\Pi, \cdot}(x, y)$. Indeed,

$$\beta_{\Pi, \cdot}(kx + s, ky + t) = (\alpha_{\cdot}(kx + s - i, ky + t))_{i \in \mathcal{S}'},$$

and each $\alpha_{\cdot}(kx + s - i, ky + t)$ depends on $\left(\alpha_{\cdot} \left(x + \left\lfloor \frac{s-i}{k} \right\rfloor - j, y \right) \right)_{j \in \mathcal{S}}$ only; the choice of \mathcal{S}' has been made so that $j - \left\lfloor \frac{s-i}{k} \right\rfloor \in \mathcal{S}'$. This concludes the proof of Proposition 4, as we can choose $E = M^{\llbracket 0; m-1 \rrbracket \times \mathcal{S}'}$, with $e(x, y)(j, i) = \alpha_j(x - i, y)$.

3.1 Example: Θ

In the case of Θ , Equation (22) becomes

$$\Xi_{2x+s}^{2y+t} = \sum_i \alpha_{\Theta, 0}(x - i, y) \Xi_{2i+s}^t + \alpha_{\Theta, 1}(x - i, y) \Xi_{2i+s}^{2+t}. \quad (24)$$

The first term is now elementary, but the second one has to be decomposed once more, i.e.

$$\Xi_{2i+s}^{2+t} = \Xi_{2i+s}^t + \Xi_{2i+s-1}^{1+t} + \Xi_{2i+s}^{1+t} + \Xi_{2i+s+1}^{1+t}, \quad (25)$$

which is the end of it if $t = 0$, but not if $t = 1$, where we get

$$\Xi_{2i+s}^3 = \Xi_{2i+s-1}^0 + \Xi_{2i+s}^0 + \Xi_{2i+s+1}^0 + \Xi_{2i+s-2}^1 + \Xi_{2i+s+2}^1. \quad (26)$$

Comparing this with (23) we can deduce the substitution rule of $\alpha_{\Theta, \cdot}$. It can then be written in the following way, where for convenience $\alpha_{\Theta, \cdot}$

is represented as $\begin{bmatrix} \alpha_{\Theta, 1} \\ \alpha_{\Theta, 0} \end{bmatrix}$:

$$\begin{array}{c}
\boxed{\alpha_{\Theta,\cdot}(x,y)} \\
\downarrow \\
\begin{array}{|c|c|}
\hline
\alpha_{\Theta,\cdot}(2x,2y+1) & \alpha_{\Theta,\cdot}(2x+1,2y+1) \\
\hline
\alpha_{\Theta,\cdot}(2x,2y) & \alpha_{\Theta,\cdot}(2x+1,2y) \\
\hline
\end{array} \\
\parallel \\
\begin{array}{|c|c|}
\hline
\alpha_{\Theta,0}(x,y) + \alpha_{\Theta,1}(x-1,y) + \alpha_{\Theta,1}(x+1,y) & 0 \\
\hline
\alpha_{\Theta,1}(x,y) & \alpha_{\Theta,1}(x,y) + \alpha_{\Theta,1}(x+1,y) \\
\hline
\alpha_{\Theta,1}(x,y) & \alpha_{\Theta,1}(x,y) + \alpha_{\Theta,1}(x+1,y) \\
\hline
\alpha_{\Theta,0}(x,y) + \alpha_{\Theta,1}(x,y) & 0 \\
\hline
\end{array}
\end{array}$$

If we follow exactly what has been said in the general case, we ought to consider the grouping $\llbracket -2; 3 \rrbracket$. However, this general bound is obviously too rough in the case of Θ , where we will just have to take $\llbracket -1; 2 \rrbracket$. We will represent the grouping in the form

$$\begin{array}{|c|c|c|c|}
\hline
\alpha_{\Theta,1}(x-1,y) & \alpha_{\Theta,1}(x,y) & \alpha_{\Theta,1}(x+1,y) & \alpha_{\Theta,1}(x+2,y) \\
\hline
\alpha_{\Theta,0}(x-1,y) & \alpha_{\Theta,0}(x,y) & \alpha_{\Theta,0}(x+1,y) & \alpha_{\Theta,0}(x+2,y) \\
\hline
\end{array}$$

The alphabet has thus size 256, and the substitution system is described by

$$\begin{array}{c}
\begin{array}{|c|c|c|c|}
\hline
a & b & c & d \\
\hline
e & f & g & h \\
\hline
\end{array} \\
\downarrow \\
\begin{array}{|c|c|c|c|c|c|c|c|}
\hline
0 & a+c+f & 0 & b+d+g & a+c+f & 0 & b+d+g & 0 \\
\hline
a+b & b & b+c & c & b & b+c & c & c+d \\
\hline
a+b & b & b+c & c & b & b+c & c & c+d \\
\hline
0 & b+f & 0 & c+g & b+f & 0 & c+g & 0 \\
\hline
\end{array}
\end{array}$$

Let us denote by A the matrix having a 1 in position a and 0 elsewhere, B the matrix having a 1 only in position b , and so on. For these matrices, we will denote the sum of matrices by a simple juxtaposition: AB will mean $A + B$, as the matrix multiplication has no meaning in this context.

Since $T_x^0 = \delta_{x,0} T_0^0$, the starting position, with which we describe the whole line number 0, is $\cdots \mid 0 \mid H \mid G \mid F \mid E \mid 0 \mid \cdots$. Since we

have, for instance, the rule $\boxed{F} \rightarrow \begin{array}{|c|c|} \hline B & A \\ \hline F & E \\ \hline \end{array}$, the graph derived from this substitution system is aperiodic; that means that, in whatever way A, B, C, \dots are represented, either as colored dots or as white dots, the

pattern converges, and the fractal structure is described by this matrix substitution system (see Section 1.4).

To calculate the fractal dimension of our spacetime diagram we use the transition matrix of the matrix substitution system, which contains the information about the images of all states. The line corresponding to F would contain a 1 in the rows of A , B , E , and F and zeros elsewhere. As every cell gives rise to four new cells the sum of all entries in each column of the matrix is 4. We thus deal with a sparse 256×256 matrix. The base 2 logarithm of the second largest eigenvalue of this matrix is the fractal dimension of the spacetime diagram (cf. for instance [15]). Here this gives a fractal dimension of $\log_2 \frac{3+\sqrt{17}}{2} \simeq 1.8325$, as also found in [8].

Let us note that up to this point our analysis for Θ is word for word valid for all CA in $\mathcal{M}_2(\mathbb{Z}_2[u, u^{-1}])$ of determinant 1 and trace $u^{-1} + 1 + u$. The additional information is only used for the actual colouring of the picture. In general all CA with the same minimal polynomial have the same substitution system, and in dimension 2 the minimal polynomial is entirely determined by the trace and the determinant. The fractal we get if we use the substitution system starting from $\cdots \begin{array}{|c|} \hline 0 \\ \hline \end{array} \begin{array}{|c|} \hline H \\ \hline \end{array} \begin{array}{|c|} \hline G \\ \hline \end{array} \begin{array}{|c|} \hline F \\ \hline \end{array} \begin{array}{|c|} \hline E \\ \hline \end{array} \begin{array}{|c|} \hline 0 \\ \hline \end{array} \cdots$ is shown in Figure 14.

In the case of Θ the connection between the substitution system and the colored picture is very simple; let us take $\xi = \begin{pmatrix} 1 \\ 0 \end{pmatrix}$ as the initial state. Then the state of cell x after y iterations is

$$\begin{aligned} \Theta_x^y \xi &= \sum_i (\alpha_{\Theta,0}(x-i, y) \Theta_i^0 + \alpha_{\Theta,1}(x-i, y) \Theta_i^1) \xi \\ &= \alpha_{\Theta,0}(x, y) \xi + \sum_i \alpha_{\Theta,1}(x-i, y) \Theta_i^1 \xi. \end{aligned} \quad (27)$$

Since $\Theta_0^1 = \begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix}$, $\Theta_1^1 = \Theta_{-1}^1 = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}$ and $\Theta_i^1 = 0$ when $i \notin \{-1; 0; 1\}$, we have $T_x^y \xi = \begin{pmatrix} \alpha_{\Theta,0}(x, y) \\ \alpha_{\Theta,1}(x, y) \end{pmatrix}$. This gives us a color assignment for each state of the matrix substitution system, which corresponds to simply dropping all states that include neither B nor F .

We can now determine the average hue of the spacetime diagram making use of the eigenvector corresponding to the second largest eigenvalue of the transition matrix [15]. Let us say $\begin{pmatrix} 1 \\ 0 \end{pmatrix}$, $\begin{pmatrix} 0 \\ 1 \end{pmatrix}$ and $\begin{pmatrix} 1 \\ 1 \end{pmatrix}$ are respectively coded by the colors c_{10} , c_{01} and c_{11} ; let c_q be the white color. We determine which symbols of the alphabet belong to each




Figure 14: The general fractal time evolution of a linear CA with $k = m = 2$, determinant 1 and trace $u^{-1} + 1 + u$. Only the areas where the whole group of α s is 0 is marked white. Thus the image appears to have less white than the colored picture. In the limit of infinite recursion this effect vanishes, thus the fractal that is generated is actually the same.

of the colors by looking only at the part $\binom{\alpha_{\Theta,0}(x,y)}{\alpha_{\Theta,1}(x,y)}$. Then we just add up all the weights of symbols with the same color in the eigenvector. We get the following unnormalised coefficients: c_{10} : $2(4 + \sqrt{17})$, c_{01} : $2(4 + \sqrt{17})$, and c_{11} : $5 + \sqrt{17}$.

In figure 1a, this color code was used: $c_{10} = \text{yellow}$, $c_{01} = \text{blue}$ and $c_{11} = \text{red}$. We must therefore have the following average hue: brown .

CONCLUSION

We have shown that every cellular automaton inducing a morphism of abelian groups produces a spacetime diagrams that converge to selfsimilar patterns that are superpositions of elementary patterns obtained for groups of prime power order. We exhibited an algorithm taking as input the local transition rule and outputting a description of these elementary patterns. We only studied the one-dimensional case in this article, but the analysis can be carried over to higher dimensions with not much more ado. Instead of $\mathcal{M}_d(R[u, u^{-1}])$, a n -dimensional linear cellular automaton would then be an element of $\mathcal{M}_d(R[u_1, u_1^{-1}, u_2, u_2^{-1}, \dots, u_n, u_n^{-1}])$, matrix substitution systems would

become $(n + 1)$ -dimensional array substitution systems, the system of indices in section 3 would be further complicated, and the spacetime diagrams would be harder to display. However, the generalization does not present any theoretical difficulty and makes nice pictures.

For instance, Figure 15 is a view of the the fractal structure produced by $\Omega = \begin{pmatrix} 0 & 1 \\ 1 & (1 + u + u^{-1})(1 + v + v^{-1}) \end{pmatrix}$, a variation on Θ belonging to $\mathcal{M}_2(\mathbb{Z}_2[u, u^{-1}, v, v^{-1}])$. For obscure historical reasons, pyramids used to be constructed with their points facing up, so to continue the tradition, this fractal is represented upside-down compared to the two-dimensional figures. Its neighborhood being a square centered on $(0, 0)$, the pyramid has unsurprisingly a square base. We cut it open so as to reveal its insides. The sector that is cut out is delimited by two half-lines, the one that is visible on the right along $(1, 0)$, the sector having an arbitrary angle of approximately $\frac{2}{3}\pi$.




Figure 15: A modern funeral labyrinth.

For any positive integer n , the coefficient in v^n of Ω^n , seen as an element of $\mathcal{M}_2(\mathbb{Z}_2[u, u^{-1}])[v, v^{-1}]$, is $\begin{pmatrix} 0 & 0 \\ 0 & (1 + u + u^{-1})^n \end{pmatrix}$. Therefore, the pattern on the face of the pyramids is given by $(1 + u + u^{-1}) \in \mathcal{M}_1(\mathbb{Z}_2[u, u^{-1}])$, as can be indeed noticed on the figure, where half of this symmetric pattern is clearly visible on the external face. By contrast, the coefficient of Ω^n in v^0 is exactly Θ^n , which explains why the cut along $(1, 0)$ in Figure 15 looks like the half of Figure 14 (turned

upside-down).

Here are some open questions and possible future developments:

- Presumably, the m in Theorem 1 can always be taken to be 1. This is known to be true in the cyclic case, i.e. when $d = 1$, cf [14].
- The algorithm presented in this article, producing a description of the spacetime diagram in the form of a matrix substitution system, has a high complexity, due to the large size of its output. Is this a necessary evil, or can more efficient descriptions be found? Could for instance the more elegant triangle-based substitution scheme presented in section 2 be naturally generalized?
- To what extent can the algebraic structure be weakened? Instead of the alphabet being an Abelian group, could we consider an Abelian monoid? Is it possible to get rid of commutativity and/or associativity?

ACKNOWLEDGEMENTS

We would like to thank Jean-Paul Allouche, Bruno Durand and Cris Moore for their useful feedback and bibliographical hints. We gratefully acknowledge the support of the Deutsche Forschungsgemeinschaft (Forschergruppe 635), the EU (projects CORNER and QICS), the Erwin Schrödinger Institute and the Rosa Luxemburg Foundation.

REFERENCES

- [1] Jean-Paul Allouche, Fritz von Haeseler, Heinz-Otto Peitgen, Antje Petersen, and Guentcho Skordev. (November 1997). Automaticity of double sequences generated by one-dimensional linear cellular automata. *Theoretical Computer Science*, 188(1–2):195–209.
- [2] Jean-Paul Allouche, Fritz von Haeseler, Heinz-Otto Peitgen, and Guentcho Skordev. (April 1996). Linear cellular automata, finite automata and Pascal’s triangle. *Discrete Applied Mathematics*, 66(1):1–22.
- [3] Johannes Gütschow. (March 2010). Entanglement generation of Clifford quantum cellular automata. *Applied Physics B*, 98(4):623–633.
- [4] Johannes Gütschow, Sonja Uphoff, Reinhard F. Werner, and Zoltán Zimborás. (January 2010). Time asymptotics and entanglement generation of Clifford quantum cellular automata. *Journal of Mathematical Physics*, 51(1).

- [5] Fritz von Haeseler, Heinz-Otto Peitgen, and Guentcho Skordev. (September 1993). Cellular automata, matrix substitutions and fractals. *Annals of Mathematics and Artificial Intelligence*, 8(3–4):345–362.
- [6] Fritz von Haeseler, Heinz-Otto Peitgen, and Guentcho Skordev. (2001). Self-similar structure of rescaled evolution sets of cellular automata. *International Journal of Bifurcation and Chaos*, 11(4):913–941.
- [7] Serge Lang. (1993). *Algebra*. Addison-Wesley publishing company, third edition.
- [8] Alan J. Macfarlane. (2004). Linear reversible second-order cellular automata and their first-order matrix equivalents. *Journal of Physics A: Mathematical and General*, 37:10791–10814.
- [9] Benoît B. Mandelbrot, Yuval Gefen, Amnon Aharony, and Jacques Peyrere. (1985). Fractals, their transfer matrices and their eigen-dimensional sequences. *Journal of Physics A: Mathematical and General*, 18:335–354.
- [10] Cristopher Moore. (1997). Quasi-linear cellular automata. *Physica D*, 103:100–132.
- [11] Cristopher Moore. (1998). Non-abelian cellular automata. *Physica D*, 111:27–41.
- [12] Dirk M. Schlingemann, Holger Vogts, and Reinhard F. Werner. (2008). On the structure of Clifford quantum cellular automata. *Journal of Mathematical Physics*, 49.
- [13] Satoshi Takahashi. (1990). Cellular automata and multifractals: dimension spectra linear cellular automata. *Physica D*, 45(1-3):36–48.
- [14] Satoshi Takahashi. (1992). Self-similarity of linear cellular automata. *Journal of Computer and System Sciences*, 44(1):114–140.
- [15] Stephen J. Willson. (1987). Computing fractal dimensions for additive cellular automata. *Physica D*, 24:190–206.